

Express Mailing Label No.: ER211528394US

PATENT APPLICATION

Docket No.: 1500.2.14

IBM Docket No.: TUC9-2003-0050US1

UNITED STATES PATENT APPLICATION

of

Yu-Cheng Hsu

and

Richard A. Ripberger

for

**APPARATUS SYSTEM AND METHOD FOR
DETERMINISTICALLY TRANSFERRING DATA BY REBOOTING
TO A DATA TRANSFER KERNEL**

1 **APPARATUS SYSTEM AND METHOD FOR**
2 **DETERMINISTICALLY TRANSFERRING DATA BY REBOOTING**
3 **TO A DATA TRANSFER KERNEL**

4
5 **BACKGROUND OF THE INVENTION**
6

7 **1. The Field of the Invention**

8 The invention relates to methods, devices, and systems for saving data during a power
9 failure. Specifically, the invention relates to methods, devices, and systems for rebooting a
10 computer with a kernel dedicated to saving data to storage devices.
11

12 **2. The Relevant Art**

13 Computer systems such as server farms or mainframe computers typically process
14 large amounts of data. Data normally resides on non-volatile mass storage devices such as
15 hard disk drives or tape drives. Data is typically transferred from non-volatile mass storage
16 to volatile memory within a computer system for processing. Recent transaction data may
17 also reside in the volatile memory of the computer.

18 Data residing in volatile memory is vulnerable to loss if the computer loses power.
19 The cost of the data loss may be very high. To protect against data loss, computers are often
20 provided with back-up power supplies. Back-up power supplies deliver uninterrupted power
21 that allows a computer to continue operation. Long-term back-up supplies can power a
22 computer until the computer is safely shut down or regular operating power is restored.

23 Providing long-term back-up power may be impractical in a data processing facility
24 that maintains a large number of computers. As a result, short-term back-up power supplies
25 storing substantial amounts of power are often used. Short-term back-up power supplies are
26 designed to supply power until a computer can safely be shut down. Data residing in the
27 computer's volatile memory must be transferred to non-volatile memory before short-term

1 power is exhausted. Rapid data transfers reduce the risk of permanent data loss and
2 minimize the cost of back-up power supplies by shortening the time a back-up supply must
3 provide power when shutting down a computer.

4 Even when a back-up power supply is available, power failures may disrupt
5 communications, peripherals, or ancillary systems, and prevent computer processes from
6 operating normally. In such conditions, processes are susceptible to completion delays and
7 faults that may result in failure to return control to the operating kernel. Furthermore, as
8 processes delay or stall, computer operation may become increasingly unstable, and the
9 computer's standard operating kernel may ultimately stall, resulting in lost data, even though
10 the computer was continuously powered with back-up power.

11 A data saving process that saves data during a power failure or other shutdown
12 operation typically runs under the computer's standard operating kernel. The data saving
13 process may be slowed by existing processes that were already running under the standard
14 operating kernel before the power failure occurred. The existing processes may stall before
15 the data transfer can be completed. Existing processes also take processing power and
16 communications bandwidth from the critical data saving process, and can potentially cause
17 unacceptable non-deterministic delays during the data transfer process. Even shutting down
18 the existing processes can be unpredictable if power failure conditions prevent normal
19 process termination. The delays of existing processes put the computer's ability to rapidly
20 and predictably save data during a power failure shutdown at risk.

21 What is needed is a device, system, and method of configuring a computer to save
22 data as rapidly as possible during a computer system shutdown. What is more particularly
23 needed is a device, system, and method of deterministically terminating existing processing
24 and configuring a computer and related subsystems to save data from volatile storage to non-
25 volatile storage.

SUMMARY OF THE INVENTION

The various elements of the present invention have been developed in response to the present state of the art, and in particular, in response to the problems and needs in the art that have not yet been fully solved by currently available means and methods for saving data during a computer system shutdown. Accordingly, the present invention provides an improved method, device, and system for rapidly and deterministically saving data to non-volatile storage under selected conditions such as a power failure.

In one aspect of the present invention, an apparatus for deterministically saving data from a computer is presented. The computer includes a processor and memory. The processor loads data from a storage device into memory. During normal computer operation, blocks of data in memory may be marked for transfer during a rapid, deterministic data save operation.

In one embodiment, a boot control module detects a condition such as a power failure requiring a rapid, deterministic data save procedure. To initiate the data save operation, the boot control module reboots the processor, resetting the processor to an initial state and initiating the reloading of a software kernel. All previously existing processes, including the standard operating kernel, are deterministically terminated by the reboot. The boot control module boots the processor with a special data transfer kernel instead of the standard operating kernel.

The data transfer kernel may be configured to exclusively support processes and interrupts required to complete the data save procedure. In one embodiment, the data transfer kernel configures the computer modules required to transfer data in a rapid deterministic manner. The data transfer kernel may also configure non-volatile storage devices to receive data.

In certain embodiments, the data transfer kernel loads a data transfer process. The data transfer process saves data blocks from the computer to the storage devices. In one embodiment, only marked data is saved and marked data blocks may not be overwritten.

1 Unmarked data blocks are not saved and may be overwritten. In another embodiment, only
2 unmarked data is saved. Saving only unmarked data facilitates backward compatibility with
3 processes that do not mark data.

4 In another aspect of the present invention, a method for rapidly saving data is
5 presented. The method operates on a computer system, identifying and marking data that
6 must be saved during a rapid, deterministic data save operation. The method further detects
7 one or more conditions that may require a data save operation. If a data save operation is
8 required, the method reboots the computer's processor with a data transfer kernel, and all
9 previously existing processes are terminated. The method for rapidly saving data may
10 configure the computer and external storage devices for a data save operation. The data
11 transfer kernel saves data to the storage devices. In one embodiment, the method only saves
12 marked data.

13 Various elements of the present invention are combined into a system for rapidly and
14 deterministically saving data. The system includes a processor and memory. The system
15 also includes one or more storage devices. In one embodiment, the memory is volatile and
16 the storage devices are non-volatile. During normal operation, the processor may mark
17 blocks of data for transfer during a rapid, deterministic data save operation. In the event of a
18 power failure or other condition requiring that the computer's data be saved, the contents of
19 the volatile memory are saved to the non-volatile storage device.

20 The system includes a boot control module. The boot control module detects
21 conditions such as power failure that require a data save operation from the memory to the
22 storage devices. The boot control module reboots the system processor. Rebooting the
23 processor terminates all the previously existing processes that could slow or stall the
24 computer's operation. The boot control module loads a data transfer kernel. The data
25 transfer kernel exclusively supports hardware and software processes required for a rapid,
26 deterministic data save procedure.

27

1 The data transfer kernel rapidly saves data from the memory to the storage device. In
2 certain embodiments, only marked data is saved. In one embodiment, no unneeded processes
3 or hardware interrupts are loaded or serviced by the data transfer kernel, and so unneeded
4 processes and hardware cannot interfere with the data transfer. In one embodiment, the data
5 transfer kernel shuts down the system when data transfer is complete.

6 The boot control module also detects a normal system boot. During normal
7 operation, the boot control module boots the system with the standard operating kernel. In
8 one embodiment, the standard operating kernel and the data transfer kernel are stored within
9 the boot control module.

10 The present invention facilitates rapid, deterministic data save operations from
11 volatile computer memory to non-volatile storage devices while protecting against process
12 instability or stalling. The present invention further decreases the time required to save data
13 prior to shutting down a computer system. A rapid, deterministic data save procedure
14 significantly reduces the probability of losing data. The present invention also reduces the
15 cost of backup power supplies needed to maintain computer functions during power failure
16 until the computer can be shut down. These and other features and advantages of the present
17 invention will become more fully apparent from the following description and appended
18 claims, or may be learned by the practice of the invention as set forth hereinafter.

BRIEF DESCRIPTION OF THE DRAWINGS

In order that the manner in which the advantages of the invention are obtained will be readily understood, a more particular description of the invention briefly described above will be rendered by reference to specific embodiments thereof, which are illustrated in the appended drawings. Understanding that these drawings depict only typical embodiments of the invention and are not therefore to be considered to be limiting of its scope, the invention will be described and explained with additional specificity and detail through the use of the accompanying drawings in which:

Figure 1 is a flow chart diagram illustrating a data save method of the prior art;

Figure 2 is a flow chart diagram illustrating one embodiment of a data save method of the present invention;

Figure 3 is a block diagram illustrating one embodiment of a computer system of the present invention;

Figure 4 is a block diagram illustrating one embodiment of a boot control module of the present invention;

Figure 5 is a block diagram of one embodiment of an initial memory map of the present invention; and

Figure 6 is a block diagram of one embodiment of a data transfer memory map of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

Many of the functional units described in this specification have been labeled as modules, in order to more particularly emphasize their implementation independence. For example, modules may be implemented in software for execution by various types of processors. An identified module of executable code may, for instance, comprise one or more physical or logical blocks of computer instructions that may, for instance, be organized as an object, procedure, or function. Nevertheless, the executables of an identified module need not be physically located together, but may comprise disparate instructions stored in different locations which, when joined logically together, comprise the module and achieve the stated purpose for the module. For example, a module of executable code could be a single instruction, or many instructions, and may even be distributed over several different code segments, among different programs, and across several memory devices.

Modules may also be implemented in hardware as electronic circuits comprising custom VLSI circuitry, off-the-shelf semiconductors such as logic chips, transistors, or other discrete components. A module may also be implemented in programmable hardware devices such as field programmable gate arrays, programmable array logic, programmable logic devices or the like.

Similarly, operational data may be identified and illustrated herein within modules, and may be embodied in any suitable form and organized within any suitable type of data structure. The operational data may be collected as a single data set, or may be distributed over different locations including over different storage devices, and may exist, at least partially, merely as electronic signals on a system or network.

Figure 1 is a flow chart diagram illustrating one embodiment of a prior art data save method 100. The method 100 includes a standard kernel processes step 110, a data save test 120, a transfer data step 130, a kernel interrupt test 140, a transfer complete test 150, a service existing processes step 160, and a shut down system step 170, and an end step 180.

1 The standard kernel processes step 110 represents normal operation of a computer,
2 including servicing one or more existing processes. The data save test 120 detects a
3 condition that requires the computer's data to be saved, such as a power failure. If a data
4 save operation is not needed, the method 100 loops to the standard kernel processes step 110.
5 If a data save operation is needed, the method 100 proceeds to the transfer data step 130.

6 The transfer data step 130 transfers a portion of data from the computer to a storage
7 device. In one embodiment, the transfer data step 130 saves data from a volatile computer
8 memory module to a non-volatile storage device. The kernel interrupt test 140 determines if
9 the standard operating kernel must service an existing process. The existing processes were
10 previously serviced during the standard kernel processes step 110.

11 If the kernel interrupt test 140 determines one or more existing processes must be
12 serviced, the method 100 loops to the transfer data step 130 via the service existing processes
13 step 160. The service existing processes step 160 services the existing processes executing
14 on the system. During the service existing processes step 160, the computer is susceptible to
15 unpredictable delays or hanging. If the kernel interrupt test 140 determines no software or
16 hardware processes must be serviced, the method 100 proceeds to the transfer complete test
17 150.

18 The transfer complete test 150 determines if all data has been transferred from the
19 computer. If some data remains to be transferred, the method 100 loops to the transfer data
20 step 130. If all data has been transferred, the method 100 proceeds to the shut down system
21 step 170. The shut down system step 170 shuts down the computer. The shut down system
22 step 170 may also shut down the storage device. The end step 180 terminates the method
23 100.

24 The data save method 100 exposes a data save operation to failure due to existing
25 processes that may stall the computer or delay completion of the data save operation. In
26 addition, terminating existing processes as part of the method 100 requires processing and
27

1 increases the risk that valuable data may be lost if the computer stalls or is powered down
2 prematurely.

3 Figure 2 is a flow chart illustrating one embodiment of a data save method 200 of the
4 present invention. The depicted data save method 200 includes a standard kernel processes
5 step 110, a data save test 120, a boot data transfer kernel step 210, a configure storage step
6 220, a load data transfer process step 230, a transfer data step 240, a transfer complete test
7 250, a shut down system step 260, and an end step 270. The method 200 terminates
8 previously existing processes that could delay a data save operation, and rapidly and
9 deterministically saves data from the computer to a storage device. Although for purposes of
10 clarity the steps of the data save method 200 are depicted in a certain sequential order,
11 execution within an actual system may be conducted in parallel and not necessarily in the
12 depicted order.

13 The standard kernel processes step 110 in one embodiment performs a computer's
14 normal operations, i.e., servicing software and hardware processes. In one embodiment, the
15 computer marks data that must be saved during a rapid, deterministic data save procedure.
16 Marking may include setting flag bits in a data line or file header. The data save test 120
17 detects a condition, such as a power failure, that requires computer data to be rapidly and
18 deterministically saved. If a data save operation is not needed, the method 200 loops to the
19 standard kernel processes step 110. If a data save operation is needed, the method 200
20 proceeds to the boot data transfer kernel step 210.

21 The boot data transfer kernel step 210 in one embodiment reboots the computer's
22 processor. Rebooting the processor deterministically terminates all existing processes.
23 Consequently, the existing processes have no opportunity to stall or run for an excessive
24 time. In addition, the boot data transfer kernel step 210 further loads a data transfer kernel.
25 The data transfer kernel is distinguished from the standard operating kernel used in normal
26 operation by supporting only a limited set of software and hardware processes that are
27 required to save data.

1 The configure storage step 220 in one embodiment configures one or more storage
2 devices to receive saved data from the computer. The load data transfer process step 230
3 loads a data transfer process running under the data transfer kernel. The data transfer process
4 is configured to rapidly save data from the computer to the storage devices with minimal
5 interruptions from other processes. The transfer data step 240 saves data from the computer
6 to the storage devices. The configure storage step 220, the load data transfer process step
7 230, and the data transfer step 240 proceed deterministically, without delays caused by other
8 processes. In one embodiment, the transfer data step 240 saves only data that was marked
9 for a rapid, deterministic data save operation. In another embodiment, the transfer data step
10 240 saves only unmarked data.

11 The transfer complete test 250 in one embodiment determines if all data has been
12 saved from the computer. In one embodiment, the transfer complete test 250 determines if
13 all marked data has been saved. If data remains to be saved, the method 200 loops to the
14 transfer data step 240. If all data has been saved, the method 200 proceeds to the shut down
15 system step 260. The shut down system step 260 shuts down the computer. In one
16 embodiment, the shut down system step 260 also shuts down the storage devices.

17 The end step 270 terminates the method 200. The method 200 reduces the risk of lost
18 data during a data save operation by terminating processes that may stall or delay computer
19 operation and by using a dedicated data transfer kernel to save data.

20 Figure 3 is a block diagram illustrating one embodiment of a computer system 300 of
21 the present invention. The computer system 300 may be used in conjunction with the method
22 200 of Figure 2 or may be conducted independent thereof. In the depicted embodiment the
23 system 300 includes a processor module 310, a memory module 320, a boot control module
24 330, an I/O module 340, and one or more storage devices 350. The system 300 supports
25 rebooting the processor module 310 to terminate existing processes and loading a data
26 transfer kernel. The data transfer kernel rapidly and deterministically saves data to the
27 storage devices 350.

1 The processor module 310 in one embodiment manipulates instructions and data
2 stored in the memory module 320. Additionally, the processor module 310 preferably further
3 instructs the I/O module 340 to store and load instructions and data to and from the storage
4 devices 350. In one embodiment, the processor module 310 marks selected data in the
5 memory module 320 for saving during a rapid, deterministic data save operation.

6 The boot control module 330 in one embodiment boots the processor module 310
7 with a standard operating kernel during a normal boot operation. The standard operating
8 kernel manages operations of the processor module 310, loading and servicing one or more
9 processes. The boot control module 330 preferably also detects conditions requiring a rapid,
10 deterministic data save operation. If a data save operation is required, the boot control
11 module 330 reboots the processor module 310, terminating all processes including the
12 standard operating kernel.

13 The boot control module 330 boots the processor module 310 with a data transfer
14 kernel rather than the standard operating kernel. In one embodiment, the data transfer kernel
15 protects marked data in the memory module 320 from being overwritten. The data transfer
16 kernel may configure the processor module 310, the memory module 320, and the I/O
17 module 340 to perform a rapid, deterministic data save operation. The data transfer kernel
18 may also configure the storage devices 350 to receive saved data.

19 The data transfer kernel may load a data transfer process onto the processor module
20 310. The data transfer process saves data from the memory module 320 to the storage devices
21 350. In one embodiment, the data transfer process saves only marked data. The data transfer
22 kernel may shut down the system 300 after saving all data.

23 The system 300 thus supports a rapid, deterministic data save operation by rebooting
24 the processor module 310 to terminate existing processes that may stall or delay the function
25 of the system 300. The boot control module 330 loads the data transfer kernel to perform the
26 data save operation without interruptions from other processes, reducing the risk of losing
27 valuable data.

1 Figure 4 is a block diagram illustrating in greater detail one embodiment of a boot
2 control module 330 of Figure 3. The boot control module 330 in the depicted embodiment
3 includes a data save detection module 410, a boot address module 420, a kernel memory
4 module 430, and a boot logic module 440. The boot control module 330 preferably reboots a
5 processor module 310 with the data transfer kernel in response to predetermined conditions
6 such as a power failure, while loading the standard operating kernel during normal
7 operational restarts.

8 The data save detection module 410 in one embodiment detects predetermined
9 conditions requiring a rapid, deterministic data save procedure. The boot logic module 440
10 in one embodiment reboots the processor module 310 if a data save operation is required.
11 Additionally, the boot logic module 440 preferably directs the boot address module 430 to
12 provide the processor module 310 with the boot address of the data transfer kernel.

13 The processor module 310 in one embodiment loads the data transfer kernel from the
14 kernel memory 430. During a normal operational restart, the boot logic module 440 directs
15 the boot address module 430 to provide the processor module 410 with the address of the
16 standard operating kernel. The processor module 310 preferably loads and executes the
17 standard operating kernel. The boot control module 330 preferably supports loading the
18 standard operating kernel during normal operational restarts while enabling a dedicated data
19 transfer kernel to be loaded under predetermined conditions to facilitate a rapid, deterministic
20 data save operation.

21 Figure 5 is a block diagram of one embodiment of an initial memory map 500 of the
22 present invention. The memory map 500 may be used in conjunction with the method 200
23 and apparatus 300 of Figures 2-4 or may be implemented independently thereof. The
24 memory map 500 includes an address space 510, a standard operating kernel 520, one or
25 more processes 530, and one or more data blocks 540. The memory map 500 illustrates the
26 location of data and processes in the memory module 320 of the computer system 300.
27

1 The address space 510 represents all available memory addresses of the memory
2 module 320. The standard kernel 520 occupies a portion of the memory module 320. One or
3 more processes 530 also reside in the memory module 320. In normal operation, the
4 standard operating kernel 520 services the processes 530. The processes 530 utilize the data
5 540 stored on the memory module 320. In one embodiment, the computer system 300 marks
6 data 540 for saving during a rapid, deterministic data save operation.

7 Figure 6 is a block diagram of one embodiment of a data transfer memory map 600 of
8 the present invention. The memory map 600 may be used in conjunction with the method
9 200 and apparatus 300 of Figures 2-4, or may be implemented independent thereof. The
10 memory map 600 includes an address space 510, a data transfer kernel 620, a data transfer
11 process 630, and one or more data blocks 540. The memory map 600 illustrates the location
12 of data 540 and processes in the memory module 320 of the computer system 300 after
13 initiation of a rapid, deterministic data save procedure.

14 In one embodiment the boot control module 320 reboots the processor module 310
15 with the data transfer kernel 620 when a rapid, deterministic data save procedure is required
16 or requested. The reboot immediately terminates the standard operating kernel 520, and the
17 existing processes 530. In one embodiment, the data transfer kernel 620 maintains the
18 integrity of the marked data 540, preventing the marked data 540 from being overwritten.

19 The data transfer kernel 620 in one embodiment configures the processor 310 and the
20 storage devices 350 to transfer data. In one embodiment, the data transfer kernel 620 loads
21 and services the data transfer process 630. The data transfer process 630 may save data 540
22 from the memory module 320 to the storage devices 350. In one embodiment, the data
23 transfer process 630 saves only marked data 540. In another embodiment, terminating
24 existing process 530 protects the computer losing data by eliminating the risk that existing
25 processes 530 will stall or delay data saving operations.

26 The present invention deterministically saves data when a computer needs to rapidly
27 shut down, such as during a system power failure. The present invention further mitigates

1 against computer instability or stalling by immediately terminating unneeded processes. The
2 present invention also decreases the time required to transfer data from a computer system in
3 order to shut down the computer system, reducing the risk of data loss and the cost of backup
4 power supplies

5 The present invention may be embodied in other specific forms without departing
6 from its spirit or essential characteristics. The described embodiments are to be considered
7 in all respects only as illustrative and not restrictive. The scope of the invention is, therefore,
8 indicated by the appended claims rather than by the foregoing description. All changes
9 which come within the meaning and range of equivalency of the claims are to be embraced
10 within their scope.

11 What is claimed is:
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27